INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-07 13:54 UTC

AI-Accelerated Exploit Window Collapse: Rising Zero-Day Exploitation and Offensive AI Capabilities Threaten Enterprise Security

SECURITY ANALYSIS | HIGH | CVSS 7.5

SCC Item ID	SCC-STY-2026-0114
Type	Security Analysis
Severity	HIGH
CVSS Base Score	7.5
Affected Products	General enterprise environments; no specific vulnerable products identified, broad organizational exposure across unpatched attack surfaces
Discovery Source	Rss:T1 Threatintel

Executive Summary

Intelligence from CrowdStrike's 2026 Global Threat Report and IBM X-Force indicates that AI is systematically compressing the window between vulnerability discovery and active exploitation, with CrowdStrike documenting an 89% year-over-year increase in AI-assisted attacks and a 42% rise in zero-days exploited before public disclosure. Traditional vulnerability management programs, built around patch cycles measured in days or weeks, are structurally misaligned with an adversary tempo now measured in hours. This shift signals a critical structural change in enterprise risk posture: detection and response strategies anchored to known-bad indicators are losing efficacy against AI-generated novel exploit variants at the precise moment offensive AI capabilities are accelerating.

Technical Analysis

The threat landscape documented across CrowdStrike's 2026 Global Threat Report and IBM X-Force intelligence reflects a decisive shift in adversary capability, not an incremental escalation. Two converging dynamics define the change.

First, the exploit window is collapsing. CrowdStrike reports a 42% increase in zero-days exploited before public disclosure, meaning vulnerability management programs that begin remediation at CVE publication are now starting the clock after exploitation has already begun. The 89% year-over-year increase in AI-assisted attacks attributed to the same report reflects the mechanism driving this compression: AI tooling allows adversaries to accelerate vulnerability analysis, proof-of-concept development, and exploit weaponization at speeds that outpace legacy patch prioritization workflows.

Second, IBM X-Force identifies a qualitative shift on the horizon: purpose-built offensive AI datasets, distinct from the incidental use of general-purpose AI tools, represent an emerging capability tier. When adversaries move from using commodity AI to training purpose-built models on exploitation data, the fidelity and novelty of AI-generated exploits will increase substantially. This capability is treated as an imminent threat, not a speculative scenario.

The MITRE ATT&CK techniques associated with this threat cluster map the full exploitation lifecycle: vulnerability scanning (T1595.002), exploit public-facing applications (T1190), client-side exploitation (T1203), exploitation of remote services (T1210), privilege escalation (T1068), and adversary capability development including custom exploit development (T1587.001) and tool acquisition (T1588.006). Phishing variants (T1566, T1566.004) and infrastructure acquisition (T1583.001) round out the pre-exploitation chain, while automated exfiltration (T1020) closes it.

Named threat actors in this context - FANCY BEAR (Russian GRU-affiliated, known for zero-day integration), FAMOUS CHOLLIMA (North Korean operator focused on software supply chain and credential theft), and PUNK SPIDER (eCrime actor with demonstrated exploitation capability) - represent different adversary tiers all trending toward AI-assisted workflows. Their presence in this intelligence cluster suggests AI-accelerated exploitation is not confined to nation-state actors; it is diffusing across the eCrime ecosystem.

The assigned CWEs (CWE-693: Protection Mechanism Failure, CWE-284: Improper Access Control, CWE-20: Improper Input Validation) reflect the vulnerability classes most susceptible to AI-assisted exploitation at scale. These are not novel weakness categories; they are the enduring structural weaknesses that AI tooling makes faster to identify and exploit.

Defensive AI adoption lagging offensive AI deployment is the central asymmetry security teams must close. Detection philosophies built around signature-based or indicator-of-compromise models are structurally disadvantaged against AI-generated novel variants that share no known-bad fingerprint with prior campaigns. These statistics are vendor-reported and should be corroborated against primary sources before use in internal risk reporting.

Action Checklist

1. Step 1: Assess exposure, audit your current mean time to patch (MTTP) for critical and high-severity vulnerabilities; determine whether your remediation SLAs assume a post-public-disclosure start, and quantify how many systems would be unpatched during a pre-disclosure exploitation window
2. Step 2: Review controls, verify EDR coverage across all internet-facing and privileged systems; confirm detection rules include behavioral indicators (T1203, T1210, T1068) not solely signature-based IOCs; review WAF and input validation controls covering CWE-20 and CWE-284 exposure classes
3. Step 3: Update threat model, incorporate AI-accelerated pre-disclosure exploitation as a named threat scenario in your threat register; map FANCY BEAR, FAMOUS CHOLLIMA, and PUNK SPIDER TTPs against your environment using MITRE ATT&CK Navigator; treat the T1595.002 to T1190 to T1068 chain as a high-priority detection engineering target
4. Step 4: Evaluate vulnerability prioritization methodology, determine whether your current program uses EPSS scoring or threat intelligence enrichment to prioritize beyond CVSS; static CVSS-based prioritization is structurally insufficient when exploitation precedes public disclosure
5. Step 5: Communicate findings, brief leadership on the patch window compression risk with specific reference to your organization's current MTTP metrics and the gap those metrics create; frame the risk in

operational terms (systems exposed during exploitation window) rather than abstract percentages

6. Step 6: Monitor developments, confirm primary source access to the CrowdStrike 2026 Global Threat Report to corroborate the 89% and 42% statistics cited; monitor IBM X-Force for follow-up publications on purpose-built offensive AI datasets; subscribe to CISA Known Exploited Vulnerabilities catalog for exploitation signal before patch cycle completion

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to CISO and legal counsel if internal log review from Step 2 (Sysmon Event ID 1, Windows Security Event ID 4688, or web server access logs) reveals T1595.002 scanning followed by T1190 exploitation attempts against internet-facing systems, or if any CVE added to the CISA KEV catalog within the past 30 days is present in your unpatched vulnerability queue on an internet-facing or privileged system — either condition indicates active exposure to the AI-accelerated pre-disclosure exploitation window described in this threat scenario.
Recovery Notes	Because this threat scenario involves pre-disclosure exploitation, recovery verification must go beyond confirming patch installation: after patching internet-facing and privileged systems, run memory forensics using Volatility3 against any system that was exposed during the MTTP window to rule out persistent implants installed before patch deployment — FANCY BEAR and FAMOUS CHOLLIMA both use persistence mechanisms (T1547, T1543) that survive patching. Monitor EDR telemetry and Sysmon logs for 30 days post-patch for re-exploitation attempts targeting the same attack surface, as AI-assisted tooling may retry against adjacent systems. Validate that WAF rules addressing CWE-20 and CWE-284 remain active and are generating logs, as adversaries may attempt input validation bypass techniques that do not require unpatched vulnerabilities.
Forensic Artifacts	Web server access logs (Apache /var/log/apache2/access.log, Nginx /var/log/nginx/access.log, IIS %SystemDrive%\inetpub\logs\LogFiles) — review for HTTP 4xx/5xx response code spikes, abnormal URI lengths, and encoded payload strings (%27, %3C, %2F sequences) that indicate CWE-20 exploitation attempts consistent with T1190 activity by PUNK SPIDER or FANCY BEAR tooling | Windows Security Event Log Event ID 4688 (Process Creation with command line auditing enabled) — filter for child processes spawned by internet-facing service executables (w3wp.exe, java.exe, tomcat.exe) including cmd.exe, powershell.exe, and wscript.exe, which indicate T1203 (Exploitation for Client Execution) or T1068 (Exploitation for Privilege Escalation) post-exploitation activity | Sysmon Event ID 3 (Network Connection) logs from internet-facing and privileged systems — identify unexpected outbound connections from service processes to external IPs following exploitation; FAMOUS CHOLLIMA tooling characteristically establishes C2 over HTTPS to non-categorized domains within minutes of initial access via T1190 | Firewall and perimeter device logs showing inbound sequential port scans or vulnerability scanner fingerprints (Shodan, Censys, custom AI-generated scanner signatures) against your internet-facing IP ranges — this is the T1595.002 (Vulnerability Scanning) precursor activity that AI-assisted tooling performs before automated exploit delivery, and its presence indicates targeting activity that predates exploitation | CISA KEV catalog diff log (daily JSON feed comparison) correlated against your open vulnerability queue export — this artifact documents the gap between when exploitation signal was available and when your organization completed remediation, which is the core forensic evidence for post-incident timeline reconstruction under NIST 800-61r3 §4 (Post-Incident Activity) and any regulatory notification review

Per-Action IR Details

Step 1: Assess exposure — audit your current mean time to patch (MTTP) for critical and high-severity vulnerabilities; determine whether your remediation SLAs assume a post-public-disclosure start, and quantify how many systems would be unpatched during a pre-disclosure exploitation window

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing IR capability and measuring organizational readiness before incidents occur

Controls: NIST SI-2 (Flaw Remediation) — requires identification, reporting, and correction of system flaws with defined timeframes, NIST RA-3 (Risk Assessment) — assess likelihood and impact of threats including pre-disclosure exploitation scenarios, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — process must account for exploitation tempo that precedes CVE publication, CIS 7.2 (Establish and Maintain a Remediation Process) — remediation SLAs must be recalibrated to reflect AI-compressed exploitation windows, not assumed post-NVD-publication start dates

Compensating: Export your patch management history from WSUS, SCCM, or Ansible Tower logs and calculate median days-to-patch per severity tier using a spreadsheet. Cross-reference against NVD publication dates using the NVD JSON feed (https://nvd.nist.gov/vuln/data-feeds) to identify the historical gap. For teams without a formal VM tool, run: 'wmic qfe list full /format:csv > patch_inventory.csv' on Windows endpoints, or 'rpm -qa --last | head -50' on Linux, and compare against CISA KEV publication dates to identify systems that were patched after exploitation was publicly known.

Evidence: Before recalibrating SLAs, preserve current-state evidence: export your ticketing system (Jira, ServiceNow) patch ticket creation timestamps versus closure timestamps for the last 12 months to establish your actual MTTP baseline. Capture WSUS/SCCM compliance reports showing unpatched systems by severity tier as of today. This establishes the pre-improvement baseline required to demonstrate risk reduction to leadership and satisfies NIST SI-2 documentation requirements.

Step 2: Review controls — verify EDR coverage across all internet-facing and privileged systems; confirm detection rules include behavioral indicators (T1203, T1210, T1068) not solely signature-based IOCs; review WAF and input validation controls covering CWE-20 and CWE-284 exposure classes

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Ensuring detection tools, coverage, and behavioral detection capability are in place before AI-accelerated zero-day exploitation occurs

Controls: NIST SI-4 (System Monitoring) — monitoring must include behavioral detection, not only signature-based methods, to address pre-disclosure zero-day exploitation, NIST SI-3 (Malicious Code Protection) — requires non-signature-based protection mechanisms; signature-only controls are structurally insufficient against AI-generated novel exploits, NIST SI-10 (Information Input Validation) — directly maps to CWE-20 (Improper Input Validation) and CWE-284 (Improper Access Control) exposure classes, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — coverage gaps in EDR and WAF must be tracked as vulnerability management findings, CIS 4.4 (Implement and Manage a Firewall on Servers) — WAF controls addressing CWE-20 and CWE-284 are a direct implementation of this safeguard for internet-facing systems

Compensating: For EDR coverage gaps: deploy Sysmon with SwiftOnSecurity's config (github.com/SwiftOnSecurity/sysmon-config) on all internet-facing and privileged Windows systems; Sysmon Event ID 1 (Process Create) and Event ID 3 (Network Connection) provide behavioral telemetry for T1203 and T1068 without commercial EDR. For T1210 (Exploitation of Remote Services): deploy osquery with the 'listening_ports' and 'process_open_sockets' queries on a 5-minute schedule to detect unexpected service behavior. For WAF coverage without budget: ModSecurity with OWASP Core Rule Set (CRS v4) addresses CWE-20 and CWE-284 at the application layer. Write Sigma rules targeting parent-child process anomalies (e.g., web server process spawning cmd.exe or powershell.exe) and run against Windows Event Log using Chainsaw (github.com/WithSecureLabs/chainsaw).

Evidence: Before completing this coverage review, capture the current EDR enrollment report showing which internet-facing and privileged systems lack coverage — this is your attack surface exposure baseline. Pull WAF rule set version and last-updated timestamps. Query Sysmon or Windows Security Event Log for Event ID 4688 (Process Creation) where ParentImage matches your internet-facing service executables (IIS w3wp.exe, Apache httpd.exe, nginx.exe) spawning interpreters (cmd.exe, powershell.exe, wscript.exe) — presence of these events before your review means behavioral detection gaps may already be producing missed signals for T1203 or T1068 activity.

Step 3: Update threat model — incorporate AI-accelerated pre-disclosure exploitation as a named threat scenario in your threat register; map FANCY BEAR, FAMOUS CHOLLIMA, and PUNK SPIDER TTPs against your environment using MITRE ATT&CK Navigator; treat the T1595.002 to T1190 to T1068 chain as a high-priority detection engineering target

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Integrating current threat intelligence into the IR capability and prioritizing detection engineering before incidents occur

Controls: NIST RA-3 (Risk Assessment) — threat model must be updated to reflect AI-accelerated adversary tempo as a named, documented threat scenario with likelihood and impact ratings, NIST IR-4 (Incident Handling) — incident handling capability must be aligned to the specific TTPs of named threat actors including FANCY BEAR (APT28), FAMOUS CHOLLIMA (Lazarus-affiliated), and PUNK SPIDER, NIST SI-4 (System Monitoring) — detection engineering for the T1595.002→T1190→T1068 chain is a direct implementation requirement of system monitoring controls, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — threat intelligence enrichment from named actor TTP profiles must feed the vulnerability prioritization process

Compensating: Load FANCY BEAR (G0007), FAMOUS CHOLLIMA (G0032 and related), and PUNK SPIDER ATT&CK group pages into MITRE ATT&CK Navigator (attack.mitre.org/resources/attack-navigator) and export the TTP heatmap as a layer file. Cross-reference each technique against your Sysmon and Windows Event Log detection coverage to identify gaps. For the T1595.002 (Vulnerability Scanning) detection: monitor your perimeter firewall or router logs for sequential port sweep patterns from single source IPs — on pfSense/OPNsense, enable Suricata with ET Open ruleset (Emerging Threats) which includes rules for automated scanner fingerprints. For T1190 (Exploit Public-Facing Application): review web server access logs for HTTP 500 response code spikes, unusual URI lengths exceeding 2,000 characters, and encoded payloads (%2F, %27, %3C) using GoAccess or grep against Apache/Nginx/IIS logs.

Evidence: Before updating the threat model, extract the last 30 days of firewall deny logs and web server error logs to establish a behavioral baseline for your environment — this allows you to determine whether T1595.002 scanning activity by FANCY BEAR or PUNK SPIDER tooling is already present before the threat model update formalizes it as a detection target. Document the current ATT&CK Navigator coverage layer as a baseline artifact so post-improvement coverage gains can be measured and reported to leadership under NIST RA-3 requirements.

Step 4: Evaluate vulnerability prioritization methodology — determine whether your current program uses EPSS scoring or threat intelligence enrichment to prioritize beyond CVSS; static CVSS-based prioritization is structurally insufficient when exploitation precedes public disclosure

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Ensuring vulnerability prioritization methodology is capable of surfacing threats that precede CVE publication, which is foundational to IR readiness in an AI-accelerated exploitation environment

Controls: NIST SI-2 (Flaw Remediation) — requires prioritization of flaw remediation; static CVSS alone does not satisfy the intent of this control when adversary exploitation precedes scoring, NIST RA-3 (Risk Assessment) — risk assessment must incorporate threat likelihood signals beyond CVSS base scores, including EPSS probability and CISA KEV membership, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — the process must be documented and reviewed; if EPSS and CTI enrichment are absent, that gap is a process deficiency requiring documented remediation, CIS 7.2 (Establish and Maintain a Remediation Process) — risk-based remediation strategy must incorporate exploitation probability signals, not solely severity scores

Compensating: Integrate EPSS scores at no cost via the FIRST EPSS API (api.first.org/epss) — query by CVE ID to retrieve the probability-of-exploitation-in-the-wild score and prioritize any CVE with EPSS > 0.10 (10%) regardless of CVSS score. Cross-reference your open vulnerability findings daily against the CISA KEV catalog (cisa.gov/known-exploited-vulnerabilities-catalog), which is freely available as a JSON feed and represents confirmed active exploitation signal that predates or supplements NVD scoring. For a 2-person team: write a simple Python or PowerShell script that pulls your scanner CSV export, queries EPSS API in bulk, appends KEV membership flag, and re-sorts by exploitation probability — this operationalizes risk-based prioritization without a commercial VM platform.

Evidence: Before changing your prioritization methodology, export your current open vulnerability queue ranked by CVSS score, then re-rank the same queue using EPSS scores from the FIRST API. The delta between the two ranked lists — specifically, low-CVSS vulnerabilities that rank high on EPSS — is forensic evidence of the structural gap that AI-accelerated exploitation exploits. Preserve both ranked lists as dated artifacts to demonstrate the methodology gap to leadership and auditors, and to establish a pre-change baseline for measuring improvement under NIST SI-2 and CIS 7.2.

Step 5: Communicate findings — brief leadership on the patch window compression risk with specific reference to your organization's current MTTP metrics and the gap those metrics create; frame the risk in operational terms (systems exposed during exploitation window) rather than abstract percentages

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing management support and organizational awareness of the IR capability gap created by AI-accelerated exploitation tempo

Controls: NIST IR-1 (Policy and Procedures) — leadership briefing is required to drive policy updates that reflect AI-accelerated threat reality; IR policy that still assumes post-disclosure patch windows is outdated and must be escalated, NIST IR-6 (Incident Reporting) — reporting to organizational leadership on systemic risk conditions (not just active incidents) is within scope of this control's intent, NIST RA-3 (Risk Assessment) — risk assessment results, including the MTTP gap against AI-compressed exploitation timelines, must be communicated to authorizing officials, CIS 7.2 (Establish and Maintain a Remediation Process) — leadership must understand and authorize the risk-based remediation strategy; this briefing is the mechanism for that authorization

Compensating: Construct the leadership briefing using three concrete data points your team already holds from Steps 1 and 4: (1) your organization's actual MTTP in days by severity tier, (2) the number of currently open critical/high CVEs where EPSS > 0.10 that remain unpatched, and (3) the count of internet-facing systems lacking EDR coverage from Step 2. Convert abstract risk into operational exposure: 'We have X internet-facing systems unpatched for an average of Y days on critical vulnerabilities; threat actors using AI-assisted tooling are weaponizing similar vulnerabilities in under 24 hours of discovery, creating a Z-day unprotected window.' No commercial tools required — this briefing is built entirely from the outputs of Steps 1-4.

Evidence: The evidence package for this briefing is the documented output of the prior four steps: the MTTP baseline report from Step 1, the EDR coverage gap report from Step 2, the ATT&CK Navigator coverage layer from Step 3, and the CVSS-vs-EPSS delta analysis from Step 4. Preserve these as dated, version-controlled artifacts — they constitute the risk assessment record required under NIST RA-3 and provide the evidentiary basis for any future audit or post-incident review questioning whether leadership was informed of the AI-accelerated exploitation risk prior to an incident.

Step 6: Monitor developments — track the CrowdStrike 2026 Global Threat Report release for primary source access to the 89% and 42% statistics cited; monitor IBM X-Force for follow-up publications on purpose-built offensive AI datasets; subscribe to CISA Known Exploited Vulnerabilities catalog for exploitation signal before patch cycle completion

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: Integrating external threat intelligence into ongoing detection and analysis capability; CISA KEV subscription provides operational exploitation signal that supplements internal detection for pre-disclosure zero-day activity

Controls: NIST SI-5 (Security Alerts, Advisories, and Directives) — requires receiving and acting on security alerts from external organizations including CISA; KEV catalog subscription is a direct implementation of this control, NIST

IR-5 (Incident Monitoring) — ongoing monitoring must incorporate external intelligence streams including CrowdStrike and IBM X-Force publications that inform the AI-accelerated threat landscape, NIST AU-6 (Audit Record Review, Analysis, and Reporting) — external intelligence publications must be correlated against internal audit records to identify whether described TTPs are present in your environment, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — the vulnerability management process must incorporate exploitation signal from CISA KEV as a continuous input, not a periodic review

Compensating: Subscribe to the CISA KEV JSON feed (cisa.gov/sites/default/files/feeds/known_exploited_vulnerabilities.json) via a free cron job or scheduled PowerShell task that pulls the feed daily, diffs against the previous day's version, and alerts your team to newly added entries — this provides exploitation signal that often precedes patch availability. For CrowdStrike and IBM X-Force primary reports: set Google Scholar and Google Alerts for 'CrowdStrike Global Threat Report 2026' and 'IBM X-Force Threat Intelligence Index 2026' to receive publication notifications without a paid subscription. Subscribe to CISA's free email alerts at cisa.gov/uscert/mailing-lists-and-feeds for advisories that reference AI-assisted exploitation campaigns targeting enterprise environments.

Evidence: Configure your KEV feed monitoring to log each newly added CVE entry with the date added, the vulnerability name, and whether that CVE is present in your open vulnerability queue — this log becomes forensic evidence of your organization's detection latency between KEV addition and internal remediation action. If an incident later occurs involving a CVE that was in the KEV catalog before your patch, this log demonstrates whether the exploitation signal was received and acted upon, which is directly relevant to NIST IR-6 (Incident Reporting) post-incident analysis and any regulatory notification timeline review.

Detection Guidance

Detection for AI-accelerated exploitation must shift from indicator-matching to behavioral baselining. Specific hunt targets derived from the MITRE techniques in this cluster:

Pre-exploitation (T1595.002): Elevated scanning activity against external-facing assets, particularly systematic port sweeps or application fingerprinting traffic. Baseline normal scanning noise against your threat surface and alert on pattern changes in volume or targeting specificity.

Initial access (T1190, T1566, T1566.004): Web application logs showing exploitation attempts against known input validation weaknesses (CWE-20); monitor for request anomalies that deviate from normal application traffic patterns rather than matching known exploit signatures. For phishing variants, monitor for link-file and spearphishing attachment patterns associated with FANCY BEAR and FAMOUS CHOLLIMA campaigns.

Execution and lateral movement (T1203, T1210, T1068): Behavioral alerts on processes spawned from browser or document rendering engines; unexpected service-to-service authentication using valid credentials in unusual sequences; privilege escalation events outside change windows.

Capability development signals (T1587.001, T1588.006): These are harder to detect internally but threat intelligence feeds should be monitored for newly observed exploit tooling associated with tracked threat actors.

Exfiltration (T1020): Automated or high-volume outbound transfers to unrecognized destinations; data transfer events outside business hours or from systems not normally involved in data movement.

Log sources to prioritize: EDR telemetry on all internet-facing and privileged systems, web application firewall logs, network flow data for east-west lateral movement, and authentication logs for privilege escalation chains.

Critical gap to audit: If your detection rules require a known-bad IOC match to fire, evaluate what percentage of your alert coverage would survive against a novel, AI-generated exploit variant with no prior signature. That gap is your AI-exploitation blind spot.

Indicators of Compromise

Type	Value	Context	Confidence
TOOL	`Pending — refer to CrowdStrike 2026 Global Threat Report for published indicators`	CrowdStrike's 2026 Global Threat Report is expected to contain specific IOCs, tooling signatures, and infrastructure indicators associated with AI-assisted exploitation campaigns; source URLs provided are T3-tier blog content and do not publish raw indicator data	LOW
TOOL	`Pending — refer to IBM X-Force threat intelligence publications on offensive AI`	IBM X-Force research on purpose-built offensive AI datasets may include associated tooling or infrastructure indicators; the source URL provided is an analytical blog post without embedded IOC data	LOW

Framework Mappings

MITRE-ATTACK

- **T1203** — Exploitation for Client Execution
- **T1210** — Exploitation of Remote Services
- **T1595.002** — Vulnerability Scanning
- **T1190** — Exploit Public-Facing Application
- **T1587.001** — Malware
- **T1566.004** — Spearphishing Voice
- **T1566** — Phishing
- **T1583.001** — Domains
- **T1020** — Automated Exfiltration
- **T1068** — Exploitation for Privilege Escalation
- **T1588.006** — Vulnerabilities

NIST-800-53R5

- **SC-7** — Boundary Protection
- **SI-2** — Flaw Remediation
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **AC-6** — Least Privilege
- **CA-8** — Penetration Testing
- **RA-5** — Vulnerability Monitoring and Scanning
- **SI-7** — Software, Firmware, and Information Integrity

- **AT-2** — Literacy Training and Awareness
- **CA-7** — Continuous Monitoring
- **SI-8** — Spam Protection
- **AC-3** — Access Enforcement
- **SI-10** — Information Input Validation
- **IR-5** — Incident Monitoring

OWASP-TOP10-2021

- **A01:2021** — Broken Access Control
- **A03:2021** — Injection

CIS-V8

- **6.1** — Establish an Access Granting Process
- **6.2** — Establish an Access Revoking Process
- **16.10** — Apply Secure Design Principles in Application Architectures
- **7.3** — Perform Automated Operating System Patch Management
- **7.4** — Perform Automated Application Patch Management
- **8.2** — Collect Audit Logs

SOC2-TSC

- **CC6.1** — The entity implements logical access security software, infrastructure, and architectures over protected information assets
- **CC9.2** — Manages risks associated with vendors and business partners

HIPAA-SECURITY

- **164.312(a)(1)** — Access Control

ISO-27001-2022

- **A.8.26** — Application security requirements
- **A.8.8** — Management of technical vulnerabilities
- **A.5.34** — Privacy and protection of personal information
- **A.5.21** — Managing information security in the ICT supply chain

NIST-CSF-2

- **DE.CM-01** — Networks and network services are monitored
- **DE.AE-08** — Incidents are declared when adverse events meet the defined incident criteria

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1203	Exploitation for Client Execution	Execution
T1210	Exploitation of Remote Services	Lateral-Movement

Technique ID	Technique Name	Tactic
T1595.002	Vulnerability Scanning	Reconnaissance
T1190	Exploit Public-Facing Application	Initial-Access
T1587.001	Malware	Resource-Development
T1566.004	Spearphishing Voice	Initial-Access
T1566	Phishing	Initial-Access
T1583.001	Domains	Resource-Development
T1020	Automated Exfiltration	Exfiltration
T1068	Exploitation for Privilege Escalation	Privilege-Escalation
T1588.006	Vulnerabilities	Resource-Development

Sources

Source	URL	Tier
Blog	https://www.crowdstrike.com/en-us/blog/tune-in-future-of-ai-powered...	T3
	https://www.crowdstrike.com/en-us/blog/crowdstrike-named-leader-gar...	T3
	https://www.ibm.com/think/x-force/understanding-future-of-offensive...	T3
	https://www.crowdstrike.com/en-us/blog/announcing-threat-ai-industr...	T3
CrowdStrike Introduces Falcon Data Security to Stop Data Theft	https://www.crowdstrike.com/en-us/press-releases/crowdstrike-introd...	T3

Generated 2026-05-07 13:54 UTC by TJS Security Command Center